

Robust optimal quantum learning without quantum memory

G. Sentís,¹ J. Calsamiglia,¹ R. Muñoz-Tapia,¹ and E. Bagan^{1,2,3}

¹*Física Teòrica: Informació i Fenòmens Quàntics,*

Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain

²*Department of Physics, Hunter College of the City University of New York, 695 Park Avenue, New York, NY 10021, USA*

³*Physics Department, Brookhaven National Laboratory, Upton, NY 11973, USA*

A quantum learning machine for binary classification of qubit states that does not require quantum memory is introduced and shown to perform with the minimum error rate allowed by quantum mechanics for *any* size of the training set. This result is shown to be robust under (an arbitrary amount of) noise and under (statistical) variations in the composition of the training set, provided it is large enough. This machine can be used an arbitrary number of times without retraining. Its required classical memory grows only logarithmically with the number of training qubits, while its excess risk decreases as the inverse of this number, and twice as fast as the excess risk of an “estimate-and-discriminate” machine, which estimates the states of the training qubits and classifies the data qubit with a discrimination protocol tailored to the obtained estimates.

Quantum computers are expected to perform some (classical) computational tasks of practical interest, e.g., large integer factorization, with unprecedented efficiency. Quantum simulators, on the other hand, perform tasks of a more “quantum nature”, which cannot be efficiently carried out by a classical computer. Namely, they have the ability to simulate complex quantum dynamical systems of interest. The need to perform tasks of genuine quantum nature is emerging as individual quantum systems play a more prominent role in labs (and, eventually, in everyday life). Examples include: quantum teleportation, dynamical control of quantum systems, or quantum state identification. Quantum information techniques are already being developed in order to execute these tasks efficiently.

This paper is concerned with a simple, yet fundamental instance of quantum state identification. A source produces two unknown pure qubit states with equal probability. A human expert (who knows the source specifications, for instance) classifies a number of $2n$ states produced by this source into two sets of size roughly n (statistical fluctuations of order \sqrt{n} should be expected) and attaches the labels 0 and 1 to them. We view these $2n$ states as a training sample, and we set ourselves to find a universal machine that uses this sample to assign the right label to a new unknown state produced by the same source. We refer to this task as quantum classification for short.

Quantum classification can be understood as a *supervised quantum learning* problem, as has been noticed by Guta and Kotlowski in their recent work [1] (though they use a slightly different setting). Learning theory, more properly named *machine learning* theory, is a very active and broad field which roughly speaking deals with algorithms capable of learning from experience [2]. Its quantum counterpart [3–7] not only provides improvements over some classical learning problems but also has a wider range of applicability, which includes the problem at hand. Quantum learning has also strong links with quantum control theory and is becoming a significant element of the quantum information processing toolbox.

An absolute limit on the minimum error in quantum classification is provided by the so called optimal programmable discrimination machine [8–11]. In this context, to ensure optimality one assumes that a fully general two-outcome joint measurement is performed on *both* the $2n$ training qubits and the qubit we wish to classify, where the observed outcome determines which of the two labels, 0 or 1, is assigned to the latter qubit. Thus, in principle, this assumption implies that in a learning scenario a quantum memory is needed to store the training sample till the very moment we wish to classify the unknown qubit. The issue of whether or not the joint measurement assumption can be relaxed has not yet been addressed. Nor has the issue of how the information left after the joint measurement can be used to classify a second unknown qubit produced by the same source, unless a fresh new training set (TS) is provided (which may seem unnatural in a learning context).

The aim of this paper is to show that for a sizable TS (asymptotically large n) the lower bound on the probability of misclassifying the unknown qubit set by programmable discrimination can be attained by first performing a suitable measurement on the TS followed by a Stern-Gerlach type of measurement on the unknown qubit, where forward classical communication is used to control the parameters of the second measurement. The whole protocol can thus be understood as a learning machine (LM), which requires much less demanding assumptions while still having the same accuracy as the optimal programmable discrimination machine. All the relevant information about the TS needed to control the Stern-Gerlach measurement is kept in a *classical* memory, thus classification can be executed any time after the learning process is completed. Once trained, this machine can be subsequently used an arbitrary number of times to classify states produced by the same source. Moreover, this optimal LM is robust under noise, i.e., it still attains optimal performance if the states produced by the source undergo depolarization to any degree. Interestingly enough, in the ideal scenario where the qubit states are pure and the TS consists in exactly the same

number of copies of each of the two types 0/1 (no statistical fluctuations are allowed) this LM attains the optimal programmable discrimination bound for *any* size $2n$ of the TS, not necessarily asymptotically large.

At this point it should be noted that LMs without quantum memory can be naturally assembled from two quantum information primitives: state estimation and state discrimination. We will refer to these specific constructions as “estimate-and-discriminate” (E&D) machines. The protocol they execute is as follows: by performing, e.g., an optimal covariant measurement on the n qubits in the TS labeled 0, their state $|\psi_0\rangle$ is estimated with some accuracy, and likewise the state $|\psi_1\rangle$ of the other n qubits that carry the label 1 is characterized. This classical information is stored and subsequently used to discriminate an unknown qubit state. It will be shown that the excess risk (i.e., excess average error over classification when the states $|\psi_0\rangle$ and $|\psi_1\rangle$ are perfectly known) of this protocol is twice that of the optimal LM. The fact that the E&D machine is suboptimal means that the kind of information retrieved from the TS and stored in the classical memory of the optimal LM is specific to the classification problem at hand, and that the machine itself is more than the mere assemblage of well known protocols.

We will first present our results for the ideal scenario where states are pure and no statistical fluctuation in the number of copies of each type of state is allowed. The effect of these fluctuations and the robustness of the LM optimality against noise will be postponed to the end of the section.

RESULTS

Programmable machines. Before presenting our results, let us summarize what is known about optimal machines for programmable discrimination. This will also allow us to introduce our notation and conventions. Neglecting statistical fluctuations, the TS of size $2n$ is given by a state pattern of the form $[\psi_0^{\otimes n}] \otimes [\psi_1^{\otimes n}]$, where the shorthand notation $[\cdot] \equiv |\cdot\rangle\langle\cdot|$ will be used throughout the paper, and where no knowledge about the actual states $|\psi_0\rangle$ and $|\psi_1\rangle$ is assumed (the figure of merit will be an average over *all* states of this form). The qubit state that we wish to label (the *data* qubit) belongs either to the first group (it is $[\psi_0]$) or to the second one (it is $[\psi_1]$). Thus, the optimal machine must discriminate between the two possible states: either $\varrho_0^n = [\psi_0^{\otimes(n+1)}]_{AB} \otimes [\psi_1^{\otimes n}]_C$, in which case it should output the label 0, or $\varrho_1^n = [\psi_0^{\otimes n}]_A \otimes [\psi_1^{\otimes(n+1)}]_{BC}$, in which case the machine should output the label 1. Here and when needed for clarity, we name the three subsystems involved in this problem A , B and C , where AC is the TS and B is the data qubit. In order to discriminate ϱ_0^n from ϱ_1^n , a joined two-outcome measurement, independent of the actual states $|\psi_0\rangle$ and $|\psi_1\rangle$, is performed on all $2n+1$ qubits. Mathematically, it is represented by a positive operator

valued measure (POVM) $\mathcal{E} = \{E_0, E_1 = \mathbb{1} - E_0\}$. The minimum average error probability of the quantum classification process is given by $P_e = (1 - \Delta/2)/2$, where $\Delta = 2 \max_{E_0} \int d\psi_0 d\psi_1 \text{tr}[(\varrho_0^n - \varrho_1^n) E_0]$. This average can be cast as a $SU(2)$ group integral and, in turn, readily computed using Schur’s lemma to give

$$\Delta = 2 \max_{E_0} \text{tr}[(\sigma_0^n - \sigma_1^n) E_0] = \|\sigma_0^n - \sigma_1^n\|_1, \quad (1)$$

where $\|\cdot\|_1$ is the trace norm and $\sigma_{0/1}^n$ are average states defined as

$$\begin{aligned} \sigma_0^n &= \frac{\mathbb{1}_{n+1} \otimes \mathbb{1}_n}{d_{n+1} d_n} = \frac{\mathbb{1}_{AB} \otimes \mathbb{1}_C}{d_{AB} d_C}, \\ \sigma_1^n &= \frac{\mathbb{1}_n \otimes \mathbb{1}_{n+1}}{d_n d_{n+1}} = \frac{\mathbb{1}_A \otimes \mathbb{1}_{BC}}{d_A d_{BC}}. \end{aligned} \quad (2)$$

In this paper $\mathbb{1}_m$ stands for the projector on the fully symmetric invariant subspace of m qubits, which has dimension $d_m = m+1$. Sometimes, it turns out to be more convenient to use the subsystem labels, as on the right of (2). The maximum in (1) is attained by choosing E_0 to be the projector onto the positive part of $\sigma_0^n - \sigma_1^n$.

The right-hand side of (1) can be computed by switching to the total angular momentum basis, $\{|J, M\rangle\}$, where $\frac{1}{2} \leq J \leq n + \frac{1}{2}$ and $-J \leq M \leq J$ (an additional label may be required to specify the way subsystems couple to give J ; see below). In this (Jordan) basis [10] the problem simplifies significantly, as it reduces to pure state discrimination [12] on each subspace corresponding to a possible value of the total angular momentum J and magnetic number M . By writing the various values of the total angular momentum as $J = k + \frac{1}{2}$, the final answer takes the form [11]:

$$P_e^{\text{opt}} = \frac{1}{2} - \frac{1}{d_n^2 d_{n+1}} \sum_{k=0}^n k \sqrt{d_n^2 - k^2}. \quad (3)$$

A simple asymptotic expression for large n can be computed using Euler-Maclaurin’s summation formula. After some algebra one obtains

$$P_e^{\text{opt}} \simeq \frac{1}{6} + \frac{1}{3n}. \quad (4)$$

The leading order (1/6) coincides with the average error probability $\int d\psi_0 d\psi_1 p_e^{\text{opt}}(\psi_0, \psi_1)$, where $p_e^{\text{opt}}(\psi_0, \psi_1)$ is the minimum error in discrimination between the two *known* states $|\psi_0\rangle$ and $|\psi_1\rangle$.

Learning machines. The formulas above give an absolute lower bound to the error probability that can be physically attainable. We wish to show that this bound can actually be attained by a learning machine that uses a classical register to store all the relevant information obtained in the learning process regardless the size, $2n$, of the TS. A first hint that this may be possible is that the optimal measurement \mathcal{E} can be shown to have positive partial transposition with respect to the partition

TS/data qubit. Indeed this is a necessary condition for any measurement that consists of a local POVM on the TS whose outcome is fed-forward to a second POVM on the data qubit. This class of one-way adaptive measurement can be characterized as:

$$E_0 = \sum_{\mu} L_{\mu} \otimes D_{\mu}, \quad E_1 = \sum_{\mu} L_{\mu} \otimes (\mathbb{1}_1 - D_{\mu}), \quad (5)$$

where the positive operators L_{μ} (D_{μ}) act on the Hilbert space of the TS (data qubit we wish to classify), and $\sum_{\mu} L_{\mu} = \mathbb{1}_n \otimes \mathbb{1}_n$. The POVM $\mathcal{L} = \{L_{\mu}\}$ represents the learning process, and the parameter μ , which a priori may be discrete or continuous, encodes the information gathered in the measurement and required at the classification stage. For each possible value of μ , $\mathcal{D}_{\mu} = \{D_{\mu}, \mathbb{1}_1 - D_{\mu}\}$ defines the measurement on the data qubit, whose two outcomes represent the classification decision. Clearly, the size of the required classical memory will be determined by the information content of the random variable μ .

Covariance and structure of \mathcal{L} . We will next prove that the POVM \mathcal{L} , which extracts the relevant information from the TS, can be chosen to be covariant. This will also shed some light on the physical interpretation of the classical variable μ . The states (2) are by definition invariant under a rigid rotation acting on subsystems AC and B , of the form $U = U_{AC} \otimes u$, where throughout this paper, U stands for an element of the appropriate representation of $SU(2)$, which should be obvious by context (in this case $U_{AC} = u^{\otimes 2n}$, where u is in the fundamental representation). Since $\text{tr}(E_0 \sigma_{0/1}^n) = \text{tr}(E_0 U^{\dagger} \sigma_{0/1}^n U) = \text{tr}(U E_0 U^{\dagger} \sigma_{0/1}^n)$, the positive operator $U E_0 U^{\dagger}$ gives the same error probability as E_0 for *any* choice of U [as can be seen from, e.g., Eq. (1)]. The same property thus holds for their average over the whole $SU(2)$ group $\bar{E}_0 = \int du U E_0 U^{\dagger}$, which is invariant under rotations, and where du denotes the $SU(2)$ Haar measure. By further exploiting rotation invariance (see Sec. Methods for full details) \bar{E}_0 can be written as

$$\bar{E}_0 = \int du \left(U_{AC} \Omega U_{AC}^{\dagger} \right) \otimes (u[\uparrow]u^{\dagger}) \quad (6)$$

for some positive operator Ω , where we use the short hand notation $[\uparrow] \equiv |\frac{1}{2}, \frac{1}{2}\rangle\langle\frac{1}{2}, \frac{1}{2}|$. Similarly, the second POVM element can be chosen to be an average, \bar{E}_1 , of the form (6), with $[\downarrow] \equiv |\frac{1}{2}, -\frac{1}{2}\rangle\langle\frac{1}{2}, -\frac{1}{2}|$ instead of $[\uparrow]$. We immediately recognize $\bar{\mathcal{E}} = \{\bar{E}_0, \bar{E}_1\}$ to be of the form (5), where u , $L_u \equiv U_{AC} \Omega U_{AC}^{\dagger}$ and $D_u \equiv u[\uparrow]u^{\dagger}$ play the role of μ , L_{μ} and D_{μ} respectively. Hence, w.l.o.g. we can choose $\mathcal{L} = \{U_{AC} \Omega U_{AC}^{\dagger}\}_{SU(2)}$, which is a covariant POVM with seed Ω . Note that u entirely defines the Stern-Gerlach measurement, $\mathcal{D}_u = \{u[\uparrow]u^{\dagger}, u[\downarrow]u^{\dagger}\}$, i.e., u specifies the direction along which the Stern-Gerlach has to be oriented. This is the relevant information that has to be retrieved from the TS and kept in the classical memory of the LM.

Covariance has also implications on the structure of Ω . In Sec. Methods, we show that this seed can always be written as

$$\Omega = \sum_{m=-n}^n \Omega_m; \quad \Omega_m \geq 0, \quad (7)$$

where

$$\sum_{m=-j}^j \langle j, m | \Omega_m | j, m \rangle = 2j + 1, \quad 0 \leq j \leq n, \quad (8)$$

and j (m) stands for the total angular momentum j_{AC} (magnetic number m_{AC}) of the qubits in the TS. In other words, the seed is a direct sum of operators with a well defined magnetic number. As a result, we can interpret that Ω points along the z -axis. The constrain (8) ensures that \mathcal{L} is a resolution of the identity.

To gain more insight into the structure of Ω , we trace subsystems B in the definition of Δ , given by the first equality in Eq. (1). For the covariant POVM (6), rotational invariance enables us to express this quantity as

$$\Delta^{LM} = 2 \max_{\Omega} \text{tr}\{(\sigma_0^n - \sigma_1^n) \Omega \otimes [\uparrow]\} = 2 \max_{\Omega} \text{tr}(\Gamma_{\uparrow} \Omega), \quad (9)$$

where we have defined

$$\Gamma_{\uparrow} = \text{tr}_B\{[\uparrow](\sigma_0^n - \sigma_1^n)\} \quad (10)$$

(the two resulting terms in the right-hand side are the post-measurement terms of AC conditioned to the outcome \uparrow after the Stern-Gerlach measurement \mathcal{D}_z is performed on B) and the maximization is over valid seeds (i.e., over positive operators Ω such that $\int du U_{AC} \Omega U_{AC}^{\dagger} = \mathbb{1}_{AC}$). We calculate Γ_{\uparrow} in Sec. Methods. The resulting expression can be cast in the simple and transparent form

$$\Gamma_{\uparrow} = \frac{\hat{J}_z^A - \hat{J}_z^C}{d_n^2 d_{n+1}}, \quad (11)$$

where $\hat{J}_z^{A/C}$ is the z component of the total angular momentum operator acting on subsystem A/C , i.e., on the training qubits to which the human expert assigned the label 0/1. Eq. (11) suggests that the optimal Ω should project on the subspace of A (C) with maximum (minimum) magnetic number, which implies that $m_{AC} = 0$. An obvious candidate is

$$\Omega = [\phi^0], \quad |\phi^0\rangle = \sum_{j=0}^n \sqrt{2j+1} |j, 0\rangle. \quad (12)$$

Below we prove that indeed this seed generates the optimal LM POVM.

Optimality of the LM. We now prove our main result: the POVM $\bar{\mathcal{E}} = \{\bar{E}_0, \bar{E}_1\}$, generated from the seed state in Eq. (12), gives an error probability $P_e^{LM} = (1 -$

$\Delta^{\text{LM}}/2$) equal to the minimum error probability P_e^{opt} of the optimal programmable discriminator, Eq. (3). It is, therefore, optimal and, moreover, it attains the absolute minimum allowed by quantum physics.

The proof goes as follows. From the very definition of error probability,

$$P_e^{\text{LM}} = \frac{1}{2} (\text{tr } \sigma_1^n \bar{E}_0 + \text{tr } \sigma_0^n \bar{E}_1), \quad (13)$$

we have

$$P_e^{\text{LM}} = \frac{\text{tr}(\mathbb{1}_A \otimes \mathbb{1}_{BC} [\phi^0] \otimes |\uparrow\rangle\langle\uparrow|) + \text{tr}(\mathbb{1}_{AB} \otimes \mathbb{1}_C [\phi^0] \otimes |\downarrow\rangle\langle\downarrow|)}{2d_n d_{n+1}}, \quad (14)$$

where we have used rotational invariance. We can further simplify this expression by writing it as

$$P_e^{\text{LM}} = \frac{\|\mathbb{1}_A \otimes \mathbb{1}_{BC} |\phi^0\rangle |\uparrow\rangle\|^2 + \|\mathbb{1}_{AB} \otimes \mathbb{1}_C |\phi^0\rangle |\downarrow\rangle\|^2}{2d_n d_{n+1}}. \quad (15)$$

To compute the projections inside the norm signs we first write $|\phi^0\rangle |\uparrow\rangle$ ($|\phi^0\rangle |\downarrow\rangle$) will be considered below) in the total angular momentum basis $|J, M\rangle_{(AC)B}$, where the attached subscripts remind us how subsystems A , B and C are both *ordered* and *coupled* to give the total angular momentum J (note that a permutation of subsystems, prior to fixing the coupling, can only give rise to a global phase, thus not affecting the value of the norm we wish to compute). This is a trivial task since $|\phi^0\rangle |\uparrow\rangle \equiv |\phi^0\rangle_{AC} |\uparrow\rangle_B$, i. e., subsystems are ordered and coupled as the subscript $(AC)B$ specifies, so we just need the Clebsch-Gordan coefficients

$$\langle j \pm \frac{1}{2}, \frac{1}{2}; j, 0; \frac{1}{2}, \frac{1}{2} \rangle = \pm \sqrt{\frac{j + \frac{1}{2} \pm \frac{1}{2}}{2j + 1}}. \quad (16)$$

The projector $\mathbb{1}_A \otimes \mathbb{1}_{BC}$, however, is naturally written as $\mathbb{1}_A \otimes \mathbb{1}_{BC} = \sum_{J,M} |J, M\rangle_{(AC)B} \langle J, M|$. This basis differs from that above in the coupling of the subsystems. To compute the projection $\mathbb{1}_A \otimes \mathbb{1}_{BC} |\phi^0\rangle |\uparrow\rangle$ we only need to know the overlaps between the two bases $_{(AC)B} \langle J, M | J, M \rangle_{(AC)B}$. Wigner's 6j-symbols provide this information as a function of the angular momenta of the various subsystems (the overlaps are computed explicitly in Sec. Methods).

Using the Clebsch-Gordan coefficients and the overlaps between the two bases, it is not difficult to obtain

$$\mathbb{1}_A \otimes \mathbb{1}_{BC} |\phi^0\rangle |\uparrow\rangle = \sum_{j=1}^{n+1} \sqrt{j} \frac{\sqrt{d_n + j} - \sqrt{d_n - j}}{\sqrt{2} d_n} |j - \frac{1}{2}, \frac{1}{2}\rangle_{(AC)B}, \quad (17)$$

An identical expression can be obtained for $\mathbb{1}_{AB} \otimes \mathbb{1}_C |\phi^0\rangle |\downarrow\rangle$ in the basis $|J, M\rangle_{(BA)C}$. To finish the proof, we compute the norm squared of (17) and substitute in (15). It is easy to check that this gives the expression of the error probability in (3), i.e., $P_e^{\text{LM}} = P_e^{\text{opt}}$.

Memory of the LM. Let us go back to the POVM condition, specifically to the minimum number of unitary transformations needed to ensure that, given a suitable discretization $\int du \rightarrow \sum_\mu p_\mu$ of (6), $\{p_\mu U_\mu [\phi^0] U_\mu^\dagger\}$ is a resolution of the identity for arbitrary n . This issue is addressed in [13], where an explicit algorithm for constructing finite POVMs, including the ones we need here, is given. From the results there, we can bound the minimum number of outcomes of \mathcal{L} by $2(n+1)(2n+1)$. This figure is important because its binary logarithm gives an upper bound to the minimum memory required. We see that it grows at most logarithmically with the size of the TS.

E&D machines. E&D machines can be discussed within this very framework, as they are particular instances of LMs. In this case the POVM \mathcal{L} has the form $L_{\alpha i} = M_\alpha \otimes M'_i$, where $\mathcal{M} = \{M_\alpha\}$ and $\mathcal{M}' = \{M'_i\}$ are themselves POVMs on the TS subsystems A and C respectively. The role of \mathcal{M} and \mathcal{M}' is to estimate (optimally) the qubit states in these subsystems [14]. The measurement on B (the data qubit) now depends on the pair of outcomes of \mathcal{M} and \mathcal{M}' : $\mathcal{D}_{\alpha i} = \{D_{\alpha i}, \mathbb{1}_1 - D_{\alpha i}\}$. It performs standard one-qubit discrimination according to the two pure-state specifications, say, the unit Bloch vectors \mathbf{s}_0^α and \mathbf{s}_1^i , estimated with \mathcal{M} and \mathcal{M}' . In this section, we wish to show that E&D machines perform worse than the optimal LM.

We start by tracing subsystems AC in Eq. (1), which for E&D reads

$$\Delta^{\text{E\&D}} = 2 \max_{\mathcal{M}, \mathcal{M}'} \text{tr}_B \max_{\{\mathcal{D}_{\alpha i}\}} \text{tr}_{AC} [(\sigma_0^n - \sigma_1^n) E_0]. \quad (18)$$

If we write $\Delta^{\text{E\&D}} = \max_{\mathcal{M}, \mathcal{M}'} \Delta_{\mathcal{M}, \mathcal{M}'}$, we have

$$\Delta_{\mathcal{M}, \mathcal{M}'} = \sum_{\alpha i} p_\alpha p'_i |\mathbf{r}_0^\alpha - \mathbf{r}_1^i|, \quad (19)$$

where \mathbf{r}_0^α and \mathbf{r}_1^i are the Bloch vectors of the data qubit states

$$\rho_\alpha^\alpha = \frac{1}{p_\alpha} \text{tr}_A \left(\frac{\mathbb{1}_{n+1}^{AB}}{d_{n+1}} M_\alpha \right), \quad \rho_1^i = \frac{1}{p'_i} \text{tr}_C \left(\frac{\mathbb{1}_{n+1}^{BC}}{d_{n+1}} M'_i \right), \quad (20)$$

conditioned to the outcomes α and i respectively, and $p_\alpha = d_n^{-1} \text{tr } M_\alpha$, $p'_i = d_n^{-1} \text{tr } M'_i$ are their probabilities. We now recall that optimal estimation necessarily requires that all elements of \mathcal{M} must be of the form $M_\alpha = c_\alpha U_\alpha [\psi^0] U_\alpha^\dagger$, where $|\psi^0\rangle = |\frac{n}{2}, \frac{n}{2}\rangle$, $c_\alpha > 0$, and $\{U_\alpha\}$ are appropriate $\text{SU}(2)$ rotations (analogous necessary conditions are required for \mathcal{M}') [15]. Substituting in Eq. (20) we obtain $p_\alpha = c_\alpha/d_n$, and

$$u_\alpha^\dagger \rho_0^\alpha u_\alpha = \frac{1}{d_{n+1}} (d_n [\uparrow] + [\downarrow]) \quad (21)$$

(a similar expression holds for ρ_1^i). This means that the Bloch vector of the data qubit conditioned to outcome α is proportional to \mathbf{s}_0^α (the Bloch vector of the corresponding estimate) and is shrunk by a factor $n/d_{n+1} =$

$n/(n+2) = \eta$. Note in passing that the shrinking factor η is independent of the measurements, provided it is optimal.

Surprisingly at first sight, POVMs that are optimal, and thus equivalent, for estimation may lead to different minimum error probabilities. In particular, the continuous covariant POVM is outperformed in the problem at hand by those with a finite number of outcomes. Optimal POVMs with few outcomes enforce large angles between the estimates \mathbf{s}_0^α and \mathbf{s}_1^i , and thus between \mathbf{r}_0^α and \mathbf{r}_1^i ($\pi/2$ in the $n = 1$ example below). This translates into increased discrimination efficiency, as shown by (19), without compromising the quality of the estimation itself. Hence the orientation of \mathcal{M} relative to \mathcal{M}' (which for two continuous POVMs does not even make sense) plays an important role, as it does the actual number of outcomes. With an increasing size of the TS, the optimal estimation POVMs require also a larger number of outcomes and the angle between the estimates decreases in average, since they tend to fill the 2-sphere isotropically. Hence, the minimum error probability is expected to approach that of two continuous POVMs. This is supported by numerical calculations. The problem of finding the optimal E&D machine for arbitrary n appears to be a hard one and is currently under investigation. Here we will give the absolute optimal E&D machine for $n = 1$ and, also, we will compute the minimum error probability for both \mathcal{M} and \mathcal{M}' being the continuous POVM that is optimal for estimation. The later, as mentioned, is expected to attain the optimal E&D error probability asymptotically.

We can obtain an upper bound on (19) by applying the Schwarz inequality. We readily find that

$$\begin{aligned} \Delta_{\mathcal{M}, \mathcal{M}'} &\leq \sqrt{\sum_{\alpha i} p_\alpha p'_i |\mathbf{r}_0^\alpha - \mathbf{r}_1^i|^2} \\ &= \sqrt{\sum_\alpha p_\alpha |\mathbf{r}_0^\alpha|^2 + \sum_i p'_i |\mathbf{r}_1^i|^2}, \end{aligned} \quad (22)$$

where we have used that $\sum_\alpha p_\alpha \mathbf{r}_0^\alpha = \sum_i p'_i \mathbf{r}_1^i = 0$, as follows from the POVM condition on \mathcal{M} and \mathcal{M}' . The maximum norm of \mathbf{r}_0^α and \mathbf{r}_1^i is bounded by $1/3$ [the shrinking factor η for $n = 1$]. Thus

$$\Delta_{\mathcal{M}, \mathcal{M}'} \leq \sqrt{2}/3 < 1/\sqrt{3} = \Delta^{\text{LM}}, \quad (23)$$

where the value of Δ^{LM} can be read off from Eq. (3). The E&D bound $\sqrt{2}/3$ is attained by the choices $M_{\uparrow/\downarrow} = [\uparrow/\downarrow]$ and $M'_{+/-} = [+/-]$, where we have used the definition $|\pm\rangle = (|\uparrow\rangle \pm |\downarrow\rangle)/\sqrt{2}$.

For arbitrary n , a simple expression for the error probability can be derived in the continuous POVM case, $\mathcal{M} = \mathcal{M}' = \{d_n U_{\mathbf{s}}[\psi^0] U_{\mathbf{s}}^\dagger\}_{\mathbf{s} \in \mathbb{S}^2}$, where \mathbf{s} is a unit vector (a point on the 2-sphere \mathbb{S}^2) and $U_{\mathbf{s}}$ is the representation of the rotation that takes the unit vector along the z -axis, \mathbf{z} , into \mathbf{s} . Here \mathbf{s} labels the outcomes of the measurement and thus plays the role of α and i . The

continuous version of (19) can be easily computed to be

$$\Delta^{\text{E\&D}} = \eta \int d\mathbf{s} |\mathbf{z} - \mathbf{s}| = \frac{4n}{3(n+2)}. \quad (24)$$

Asymptotically, we have $P_e^{\text{E\&D}} = 1/6 + 2/(3n) + \dots$. Therefore, the excess risk, which we recall is the difference between the average error probability of the machine under consideration and that of the optimal discrimination protocol for *known* qubit states ($1/6$), is $R^{\text{E\&D}} = 2/(3n) + \dots$. This is twice the excess risk of the optimal programmable machine and the optimal LM, which can be read off from Eq. (4):

$$R^{\text{LM}} = R^{\text{opt}} = \frac{1}{3n} + \dots \quad (25)$$

For $n = 1$, Eq. (23) leads to $R^{\text{E\&D}} = (4 - \sqrt{2})/12$. This value is already 15% larger than excess risk of the optimal LM: $R^{\text{LM}} = (4 - \sqrt{3})/12$.

Robustness of LMs. So far we have adhered to the simplifying assumptions that the two types of states produced by the source are pure and, moreover, exactly equal in number. Neither of these two assumptions is likely to hold in practice, as both, interaction with the environment, i.e., decoherence and noise, and statistical fluctuations in the numbers of states of each type, will certainly take place. Here we prove that the performance of the optimal LM is not altered by these effects in the asymptotic limit of large TS. More precisely, the excess risk of the optimal LM remains equal to that of the optimal programmable discriminator to leading order in $1/n$ when noise and statistical fluctuations are taken into account.

Let us first consider the impact of noise, which we will assume isotropic and uncorrelated. Hence, instead of producing $|\psi_{0/1}\rangle$, the source produces copies of

$$\rho_{0/1} = r[\psi_{0/1}] + (1-r)\frac{\mathbb{1}}{2}, \quad 0 < r \leq 1. \quad (26)$$

In contrast to the pure qubits case, where $[\psi_{0/1}^{\otimes n}]$ belongs to the fully symmetric invariant subspace of maximum angular momentum $j = n/2$, the state of A/C is now a full-rank matrix of the form $\rho_{0/1}^{\otimes n}$. Hence, it has projections on all the orthogonal subspaces $\mathcal{S}_j \otimes \mathbb{C}^{\nu_j^n}$, where $\mathcal{S}_j = \text{span}(\{|j, m\rangle\}_{m=-j}^j)$ and $\mathbb{C}^{\nu_j^n}$ is the multiplicity space of the representation with total angular momentum j (see Sec. Methods for a formula of the multiplicity ν_j^n), and j is in the range from 0 ($1/2$) to $n/2$ if n is even (odd). Therefore $\rho_{0/1}^{\otimes n}$ is block-diagonal in the total angular momentum eigenbasis. The multiplicity space $\mathbb{C}^{\nu_j^n}$ carries the label of the ν_j^n different equivalent representations of given j , which arise from the various ways the individual qubits can couple to produce total angular momentum j . For permutation invariant states (such as $\rho_{0/1}^{\otimes n}$), this has no physical relevance and the only

effect of $\mathbb{C}^{\nu_j^n}$ in calculations is through its dimension ν_j^n . Hence, the multiplicity space will be dropped throughout this paper.

The average states now become a direct sum of the form

$$\int d\psi_0 d\psi_1 \rho_0^{\otimes(n+1)} \otimes \rho_1^{\otimes n} = \sum_{\xi} p_{\xi}^n \sigma_{0,\xi}^n, \quad (27)$$

$$\int d\psi_0 d\psi_1 \rho_0^{\otimes n} \otimes \rho_1^{\otimes(n+1)} = \sum_{\xi} p_{\xi}^n \sigma_{1,\xi}^n, \quad (28)$$

where we use the shorthand notation $\xi = \{j_A, j_C\}$ [each angular momentum ranges from 0 (1/2) to $n/2$ for n even (odd)], and $p_{\xi}^n = p_{j_A}^n p_{j_C}^n$ is the probability of any of the two average states projecting on the block labeled ξ . Hence,

$$\Delta^{\text{LM}} = \sum_{\xi} p_{\xi}^n \|\sigma_{0,\xi}^n - \sigma_{1,\xi}^n\|_1. \quad (29)$$

The number of terms in Eq. (29), is $[(2n+3 \pm 1)/4]^2$ for even/odd n . It grows quadratically with n , in contrast to the pure state case for which there is a single contribution corresponding to $j_A = j_C = n/2$. In the asymptotic limit of large n , however, a big simplification arises because of the following results: for each ξ of the form $\xi = \{j, j\}$ ($j_A = j_C = j$), the following relation holds (see Sec. Methods)

$$\sigma_{0,\xi}^n - \sigma_{1,\xi}^n = \frac{r \langle \hat{J}_z \rangle_j}{j} (\sigma_0^{2j} - \sigma_1^{2j}), \quad (30)$$

where $\sigma_{0/1}^{2j}$ are the average states (2) for a number of $2j$ pure qubits. Here $\langle \hat{J}_z \rangle_j$ is the expectation value restricted to \mathcal{S}_j of the z -component of the angular momentum in the state $\rho^{\otimes n}$, where ρ has Bloch vector $r\mathbf{z}$. Eq. (30) is an exact algebraic identity that holds for any value of j , n and r (it bears no relation whatsoever to measurements of any kind). The second result is that for large n , both $p_{j_A}^n$ and $p_{j_C}^n$ become continuous probability distributions, $p_n(x_A)$ and $p_n(x_C)$, where $x_{A/C} = 2j_{A/C}/n \in [0, 1]$. Asymptotically, they approach Dirac delta functions peaked at $x_A = x_C = r$ (see Sec. Methods). Hence, the only relevant contribution to Δ^{LM} comes from $\xi = \{rn/2, rn/2\}$. It then follows that in the asymptotic limit

$$\sum_{\xi} p_{\xi}^n (\sigma_{0,\xi}^n - \sigma_{1,\xi}^n) \simeq \frac{2 \langle \hat{J}_z \rangle_{rn/2}}{n} (\sigma_0^{rn} - \sigma_1^{rn}). \quad (31)$$

Hence, mixed-state quantum classification using a TS of size $2n$ is equivalent to its pure-state version for a TS of size $2nr$, provided n is asymptotically large. In particular, our proof of optimality above also holds for arbitrary $r \in (0, 1]$ if the TS is sizable enough, and $R^{\text{LM}} \simeq R^{\text{opt}}$. This result is much stronger than robustness against decoherence, which only would require optimality for values of r close to unity.

From Eqs. (29) and (31) one can easily compute Δ^{LM} for arbitrary r using that [18] $\langle \hat{J}_z \rangle_j \simeq j - (1-r)/(2r)$ up to exponentially vanishing terms. The trace norm of $\sigma_0^{rn} - \sigma_1^{rn}$ can be retrieved from, e.g., Eq. (25). For rn pure qubits one has $\|\sigma_0^{rn} - \sigma_1^{rn}\|_1 \simeq (4/3)[1 - 1/(rn)]$. After some trivial algebra we obtain

$$P_e^{\text{LM}} = \frac{1}{2} - \frac{r}{3} + \frac{1}{3rn} + o(n^{-1}) \quad (32)$$

for the error probability, in agreement with the optimal programmable machine value given in [11], as claimed above. This corresponds to an excess risk of

$$R^{\text{LM}} = \frac{1}{3rn} + o(n^{-1}) = R^{\text{opt}}. \quad (33)$$

In the non-asymptotic case, the sum in Eq. (29) is not restricted to $\xi = \{j, j\}$ and the calculation of the excess risk becomes very involved. Rather than attempting to obtain an analytical result, for small training samples we have resorted to a numerical optimization. We first note that Eqs. (7) through (11) define a *semidefinite programming* optimization problem (SDP), for which very efficient numerical algorithms have been developed [17]. In this framework, one maximizes the objective function Δ^{LM} [second equality in Eq. (9)] of the SDP variables $\Omega_m \geq 0$, subject to the linear condition (8). We use this approach to compute the error probability, or equivalently, the excess risk of a LM for mixed-state quantum classification of small samples ($n \leq 5$), where no analytical expression of the optimal seed is known. For mixed states the expression of Γ_{\uparrow} and Ω_m can be found in Sec. Methods, Eqs. (40) through (42).

Our results are shown in Fig. 1, where we plot R^{LM} (shaped dots) and the lower bounds given by R^{opt} (solid lines) as a function of the purity r for up to $n = 5$. We note that the excess risk of the optimal LM is always remarkably close to the absolute minimum provided by the optimal programmable machine and in the worst case ($n = 2$) it is only 0.4% larger. For $n = 1$ we see that $R^{\text{LM}} = R^{\text{opt}}$ for any value of r . This must be the case since for a single qubit in A and C one has $j_A = j_C = 1/2$, and Eq. (30) holds.

We now turn to robustness against statistical fluctuations in the number of states of each type produced by the source. In a real scenario one has to expect that $j_A = n_A/2 \neq n_C/2 = j_C$, $n_A + n_B = 2n$. Hence, Γ_{\uparrow} has the general form (40), which gives us a hint that our choice $\Omega = \Omega_{m=0}$ may not be optimal for finite n . This has been confirmed by numerical analysis using the same SDP approach discussed above. Here, we show that the asymptotic performance (for large training samples) of the optimal LM, however, is still the same as that of the optimal programmable discriminator running under the same conditions (mixed states and statistical fluctuations in $n_{A/C}$).

Asymptotically, a real source for the problem at hand will typically produce $n_{A/C} = n \pm \delta\sqrt{n}$ mixed copies of




FIG. 1. (Color online) Excess risk R^{LM} (points) and its corresponding lower bound R^{opt} (lines), both as a function of the purity r , and for values of n ranging from 1 to 5 (from top to bottom).

each type. In Sec. Methods, it is shown that the relation (31) still holds in this case if n is large. It reads

$$\sigma_{0,\xi}^n - \sigma_{1,\xi}^n \simeq r \left(1 - \frac{1-r}{nr^2} \right) (\sigma_0^{rn} - \sigma_1^{rn}) \quad (34)$$

(δ first appears at order $n^{-3/2}$). Hence, the effect of both statistical fluctuations in $n_{A/C}$ and noise (already considered above) is independent of the machine used for quantum classification (i.e., it is the same for LM, programmable machines, E&D, ...). In particular, the relation (33), $R^{\text{LM}} = R^{\text{opt}}$, between the excess rate of the optimal LM and its absolute limit given by the optimal programmable discriminator still holds asymptotically, which proves robustness.

To illustrate this, let us consider the effect of statistical fluctuations in $n_{A/C}$ for pure states. The optimal programmable machine for arbitrary n_A, n_B and n_C was presented in [11]. The error probability for the case at hand ($n_B = 1$) is in Sec. Methods. From its asymptotic expansion when n_A and n_C are both large one readily has

$$R^{\text{opt}} = \frac{1}{6} \left(\frac{1}{n_A} + \frac{1}{n_C} \right) + \dots \quad (35)$$

We see that when $n_{A/C} = n \pm \delta\sqrt{n}$ (i.e., when statistical fluctuations in $n_{A/C}$ are taken into account) one still has $R^{\text{opt}} \simeq 1/(3n) \simeq R^{\text{LM}}$.

DISCUSSION

We have presented a *supervised* quantum learning machine that classifies a single qubit prepared in a pure but otherwise unknown state after it has been trained with a number of already classified qubits. Its performance attains the absolute bound given by the optimal programmable discrimination machine. This learning machine does not require quantum memory and can also

be reused without retraining, which may save a lot of resources. The machine has been shown to be robust against noise and statistical fluctuations in the number of states of each type produced by the source. For small sized training sets the machine is very close to optimal, attaining an excess risk that is larger than the absolute lower limit by at most 0.4%. In the absence of noise and statistical fluctuations, the machine attains optimality for *any* size of the training set.

One may rise the question of whether or not the separated measurements on the training set and data qubit can be reversed in time; in a classical scenario where, e.g., one has to identify one of two faces based on a stack of training portraits, it is obvious that, without memory limitations, the order of training and data observation can be reversed (in both cases the final decision is taken based on the very same information). We will briefly show that this is not so in the quantum world. In the reversed setting, the machine first performs a measurement \mathcal{D} , with each element of rank one $u_\mu[\uparrow]u_\mu^\dagger$, and stores the information (which of the possible outcomes is obtained) in the classical memory to control the measurement to be performed on the training set in a later time. The probability of error conditioned to one of the outcomes, say \uparrow , is given by the Helstrom formula $P_e^\uparrow = (1 - \|\Gamma_\uparrow\|_1/2)/2$, where Γ_\uparrow is defined in Eq. (10). Using Eq. (11) one has $\|\Gamma_\uparrow\|_1 = d_n^{-2} d_{n+1}^{-1} \sum_{m,m'} |m - m'| = n/[3(n+1)]$. The averaged error probability is then

$$P_e^{\text{LM}} = \frac{1}{2} \left(1 - \frac{1}{6} \frac{n}{n+1} \right). \quad (36)$$

In the limit of infinite copies we obtain $P_e^{\text{LM}} \simeq 5/12$, which is way larger than $P_e^{\text{LM}} \simeq 1/6$. The same minimum error probability of Eq. (36) can be attained by performing a Stern-Gerlach measurement on the data qubit, which requires just one bit of classical memory. This is all the classical information that we can hope to retrieve from the data qubit, in agreement with Holevo's bound [19]. This clearly limits the possibilities of a correct classification—very much in the same way as in face identification with limited memory size. In contrast, the amount of classical information “sent forward” in the optimal learning machine goes as the logarithm of the size of the training sample. This asymmetry also shows that despite the separability of the measurements, non-classical correlations between the training set and the data qubit play an important role in quantum learning.

Some relevant generalizations of this work to, e.g., higher dimensional systems and arbitrarily unbalanced training sets, remain an open problem. Another challenging problem with direct practical applications in quantum control and information processing is the extension of this work to *unsupervised* machines, where no human expert classifies the training sample.

METHODS

Block-diagonal form of $\rho^{\otimes n}$

The state $\rho^{\otimes n}$ of n identical copies of a general qubit state ρ with purity r and Bloch vector $r\mathbf{s}$, has a block diagonal form in the basis of the total angular momentum given by

$$\rho^{\otimes n} = \sum_j p_j^n \rho_j \otimes \frac{\mathbb{I}_j}{\nu_j^n}.$$

Here $j = 0(1/2), \dots, n/2$ if n is even (odd), \mathbb{I}_j is the identity in the multiplicity space $\mathbb{C}^{\nu_j^n}$, of dimension ν_j^n (the multiplicity of the representation with total angular momentum j), where

$$\nu_j^n = \binom{n}{n/2-j} \frac{2j+1}{n/2+j+1}.$$

The normalized state ρ_j , which is supported on the representation subspace $\mathcal{S}_j = \text{span}\{|j, m\rangle\}$ of dimension $2j+1 = d_{2j}$, is

$$\rho_j = U_{\mathbf{s}} \left(\sum_{m=-j}^j a_m^j |j, m\rangle \right) U_{\mathbf{s}}^\dagger,$$

where

$$a_m^j = \frac{1}{c_j} \left(\frac{1-r}{2} \right)^{j-m} \left(\frac{1+r}{2} \right)^{j+m}, \quad (37)$$

and

$$c_j = \frac{1}{r} \left\{ \left(\frac{1+r}{2} \right)^{2j+1} - \left(\frac{1-r}{2} \right)^{2j+1} \right\},$$

so that $\sum_{m=-j}^j a_m^j = 1$, and we stick to our shorthand notation $[\cdot] \equiv |\cdot\rangle\langle\cdot|$, i.e., $[j, m] \equiv |j, m\rangle\langle j, m|$. The measurement on $\rho^{\otimes n}$ defined by the set of projectors on the various subspaces \mathcal{S}_j will produce ρ_j as a posterior state with probability

$$p_j^n = \nu_j^n c_j \left(\frac{1-r^2}{4} \right)^{n/2-j}.$$

One can easily check that $\sum_j p_j^n = 1$.

In the large n limit, we can replace p_j^n for a continuous probability distribution $p_n(x)$ in $[0, 1]$, where $x = 2j/n$. Applying Stirling's approximation to p_j one obtains:

$$p_n(x) \simeq \sqrt{\frac{n}{2\pi}} \frac{1}{\sqrt{1-x^2}} \frac{x(1+r)}{r(1+x)} e^{-nH(\frac{1+x}{2} \parallel \frac{1+r}{2})},$$

where $H(s \parallel t)$ is the (binary) relative entropy

$$H(s \parallel t) = s \log \frac{s}{t} + (1-s) \log \frac{1-s}{1-t}.$$

The approximation is valid for x and r both in the open unit interval $(0, 1)$. For non-vanishing r , $p_n(x)$ becomes a Dirac delta function peaked at $x = r$, $p_\infty(x) = \delta(x-r)$, which corresponds to $j = nr/2$.

Covariance and structure of \mathcal{L}

We start with a POVM element of the form $\bar{E}_0 = \int du U E_0 U^\dagger$. Since D_μ must be a rank-one projector, it can always be written as $D_\mu = u_\mu [\uparrow] u_\mu^\dagger$ for a suitable $\text{SU}(2)$ rotation u_μ . Thus,

$$\bar{E}_0 = \sum_\mu \int du \left(U_{AC} L_\mu U_{AC}^\dagger \right) \otimes (u u_\mu [\uparrow] u_\mu^\dagger u^\dagger).$$

We next use the invariance of the Haar measure du to make the change of variable: $u u_\mu \rightarrow u'$ and, accordingly, $U_{AC} \rightarrow U'_{AC} U_\mu^\dagger$. After regrouping terms we have

$$\begin{aligned} \bar{E}_0 &= \sum_\mu \int du' \left(U'_{AC} U_\mu^\dagger L_\mu U_\mu U'_{AC}^\dagger \right) \otimes (u' [\uparrow] u'^\dagger) \\ &= \int du' \left[U'_{AC} \left(\sum_\mu U_\mu^\dagger L_\mu U_\mu \right) U'_{AC}^\dagger \right] \otimes (u' [\uparrow] u'^\dagger) \\ &= \int du \left(U_{AC} \Omega U_{AC}^\dagger \right) \otimes (u [\uparrow] u^\dagger), \end{aligned}$$

where we have defined $\Omega = \sum_\mu U_\mu^\dagger L_\mu U_\mu \geq 0$. The POVM element \bar{E}_1 is obtained by replacing $[\uparrow]$ by $[\downarrow]$ in the expressions above. From the POVM condition $\sum_\mu L_\mu = \mathbb{1}_{AC}$ it immediately follows that $\int du U_{AC} \Omega U_{AC}^\dagger = \mathbb{1}_{AC}$, where $\mathbb{1}_{AC}$ is the identity on the Hilbert space of the TS, i.e., $\mathbb{1}_{AC} = \mathbb{1}_A \otimes \mathbb{1}_C$. Therefore $\mathcal{L} = \{U_{AC} \Omega U_{AC}^\dagger\}_{\text{SU}(2)}$ is a covariant POVM. The positive operator Ω is called the seed of the covariant POVM \mathcal{L} .

Now, let $u_z(\varphi)$ be a rotation about the z -axis, which leaves $[\uparrow]$ invariant. By performing the change of variables $u \rightarrow u' u_z(\varphi)$ [and $U_{AC} \rightarrow U'_{AC} U_{zAC}(\varphi)$] in the last equation above, we readily see that Ω and $U_{zAC}(\varphi) \Omega U_{zAC}^\dagger(\varphi)$ both give the same average operator \bar{E}_0 for any $\varphi \in [0, 4\pi)$. So, its average over φ ,

$$\int_0^{4\pi} \frac{d\varphi}{4\pi} U_z(\varphi) \Omega U_z^\dagger(\varphi),$$

can be used as a seed w.l.o.g., where we have dropped the subscript AC to simplify the notation. Such a seed is by construction invariant under the group of rotations about the z -axis (just like $[\uparrow]$) and, by Schur's lemma, a direct sum of operators with well defined magnetic number. Therefore, in the total angular momentum basis for AC , we can always choose the seed of \mathcal{L} as

$$\Omega = \sum_{m=-n}^n \Omega_m; \quad \Omega_m \geq 0.$$

The constrain (8) follows from the POVM condition $\mathbb{1}_{AC} = \int du U \Omega U^\dagger$ and Schur's lemma. The result also holds if A and C have different number of copies (provided they add up to $2n$). It also holds for mixed states.

Wigner's 6j-symbols

Let us consider three angular momenta j_1, j_2, j_3 that couple to give a total J . Note that there is no unique way to carry out this coupling; we might first couple j_1 and j_2 to give a resultant j_{12} , and couple this to j_3 to give J , or alternatively, we may couple j_1 to the resultant j_{23} of coupling j_2 and j_3 . Moreover, the intermediate couplings can give in principle different values of j_{12} or j_{23} which, when coupled to j_3 or j_1 , end up giving the same value of J . All these possibilities lead to linearly independent states with the same J and M , thus they must be distinguished by specifying the intermediate angular momentum and the order of coupling. There exists a unitary transformation that maps the states obtained from the two possible orderings of the coupling; Wigner's 6j-symbols [16], denoted in the next equation by $\{ \dots \}$, provide the coefficients of this transformation:

$$\begin{aligned} & \langle (j_1 j_2) j_{12}, j_3; J, M | j_1, (j_2 j_3) j_{23}; J, M \rangle \\ &= (-1)^{j_1+j_2+j_3+J} \sqrt{(2j_{12}+1)(2j_{23}+1)} \begin{Bmatrix} j_1 & j_2 & j_{12} \\ j_3 & J & j_{23} \end{Bmatrix}. \end{aligned}$$

Note that this overlap is independent of M . For the proof of optimality of the LM, we couple subsystems A, B and C in two ways: $A(CB)$ and $(AC)B$ to produce the states $|j_A, (j_C j_B) j_{CB}; J, M\rangle$ and $|(j_A j_C) j_{AC}, j_B; J, M\rangle$, which we denote by $|J, M\rangle_{A(CB)}$ and $|J, M\rangle_{(AC)B}$ respectively for short. The various angular momenta involved are fixed to $j_A = j_C = \frac{n}{2}$, $j_B = \frac{1}{2}$, $j_{AC} = j$, $j_{CB} = \frac{n}{2} + \frac{1}{2}$, whereas $J = j \pm \frac{1}{2}$. With these values, the overlaps we need are given by

$${}_{A(CB)}\langle j \pm \frac{1}{2}, \frac{1}{2} | j \pm \frac{1}{2}, \frac{1}{2} \rangle_{(AC)B} = \sqrt{\frac{n + \frac{3}{2} \pm (j + \frac{1}{2})}{2(n+1)}}.$$

Derivation of Eqs. (30) and (34)

Let us start with the general case where $\xi = \{j, j'\}$. To obtain $\sigma_{0,\xi}^n$ we first write Eqs. (27) and (28) as the $SU(2)$ group integrals

$$\begin{aligned} \sigma_{0,\xi}^n &= \int du U_{AB} \left(\sum_{m=-j}^j a_m^j [j, m]_A \otimes \rho_0^B \right) U_{AB}^\dagger \\ &\quad \otimes \int du' U'_{C} \left(\sum_{m=-j'}^{j'} a_m^{j'} [j', m]_C \right) U_C'^\dagger, \end{aligned}$$

where a_m^j is given in Eq. (37), ρ_0^B is the mixed state ρ_0 , Eq. (26), of the qubit B . We next couple A with B (more precisely, their subspaces of angular momentum j) using the Clebsch-Gordan coefficients

$$\begin{aligned} |j + \frac{1}{2}, m + \frac{1}{2} | j, m; \frac{1}{2}, \frac{1}{2} \rangle|^2 &= \frac{j+m+1}{2j+1}, \\ |j - \frac{1}{2}, m + \frac{1}{2} | j, m; \frac{1}{2}, \frac{1}{2} \rangle|^2 &= \frac{j-m}{2j+1}. \end{aligned}$$

The resulting expressions can be easily integrated using Schur lemma. Note that the integrals of crossed terms of the form $|j, m\rangle\langle j', m|$ will vanish for all $j \neq j'$. We readily obtain

$$\sigma_{0,\xi}^n = \sum_{m=-j}^j a_m^j \left(\frac{j+1+mr}{d_{2j}} \frac{\mathbb{1}_{2j+1}^{AB}}{d_{2j+1}} + \frac{j-mr}{d_{2j}} \frac{\mathbb{1}_{2j-1}^{AB}}{d_{2j-1}} \right) \otimes \frac{\mathbb{1}_{2j'}^C}{d_{2j'}},$$

where $\mathbb{1}_{2j}$ is the projector on \mathcal{S}_j and $d_{2j} = 2j+1 = \dim \mathcal{S}_j$. The superscripts attached to the various projectors specify the subsystems to which they refer. These projectors are formally equal to those used in Eq. (2) (i.e., $\mathbb{1}_{2j}$ projects on the fully symmetric subspace of $2j$ qubits) and, hence, we stick to the same notation. Note that $\text{tr} \sigma_{0,\xi}^n = 1$, as it should be.

We can further simplify this expression by introducing $\langle \hat{J}_z \rangle_j = \sum_m m a_m^j$, i.e., the expectation value of the z -component of the total angular momentum in the state ρ_j (i.e., of $\mathbb{1}_{2j} \hat{J}_z \mathbb{1}_{2j}$ in the state $\rho_{0/1}^{\otimes n}$) for a Bloch vector $r \mathbf{z}$:

$$\sigma_{0,\xi}^n = \left(\frac{j+1+r\langle \hat{J}_z \rangle_j}{d_{2j}} \frac{\mathbb{1}_{2j+1}^{AB}}{d_{2j+1}} + \frac{j-r\langle \hat{J}_z \rangle_j}{d_{2j}} \frac{\mathbb{1}_{2j-1}^{AB}}{d_{2j-1}} \right) \otimes \frac{\mathbb{1}_{2j'}^C}{d_{2j'}}.$$

Using the relation

$$\mathbb{1}_{2j-1}^{AB} = \mathbb{1}_{2j}^A \otimes \mathbb{1}_1^B - \mathbb{1}_{2j+1}^{AB},$$

and $(j+1)/d_{2j+1} = j/d_{2j-1} = 1/2$, we can write

$$\sigma_{0,\xi}^n = \left(\frac{r\langle \hat{J}_z \rangle_j}{j} \frac{\mathbb{1}_{2j+1}^{AB}}{d_{2j+1}} + \frac{j-r\langle \hat{J}_z \rangle_j}{j} \frac{\mathbb{1}_{2j}^A}{d_{2j}} \otimes \frac{\mathbb{1}_1^B}{2} \right) \otimes \frac{\mathbb{1}_{2j'}^C}{d_{2j'}}. \quad (38)$$

Similarly, we can show that

$$\sigma_{1,\xi}^n = \frac{\mathbb{1}_{2j}^A}{d_{2j}} \otimes \left(\frac{r\langle \hat{J}_z \rangle_{j'}}{j'} \frac{\mathbb{1}_{2j'+1}^{BC}}{d_{2j'+1}} + \frac{j'-r\langle \hat{J}_z \rangle_{j'}}{j'} \frac{\mathbb{1}_1^B}{2} \otimes \frac{\mathbb{1}_{2j'}^C}{d_{2j'}} \right). \quad (39)$$

Therefore, if $j' = j$,

$$\sigma_{0,\xi}^n - \sigma_{1,\xi}^n = \frac{r\langle \hat{J}_z \rangle_j}{j} \left(\frac{\mathbb{1}_{2j+1}^{AB}}{d_{2j+1}} \otimes \frac{\mathbb{1}_{2j}^C}{d_{2j}} - \frac{\mathbb{1}_{2j}^A}{d_{2j}} \otimes \frac{\mathbb{1}_{2j+1}^{BC}}{d_{2j+1}} \right).$$

Comparing with Eq. (2), the two terms in the second line can be understood as the average states for a number of $2j$ pure qubits, i.e., as σ_0^{2j} and σ_1^{2j} respectively. Hence, if $\xi = \{j, j\}$ we have the relation

$$\sigma_{0,\xi}^n - \sigma_{1,\xi}^n = \frac{r\langle \hat{J}_z \rangle_j}{j} \left(\sigma_0^{2j} - \sigma_1^{2j} \right),$$

which is Eq. (30). It is important to emphasize that this equation is exact (i.e., it holds for any value of j, n and r) and bears no relation whatsoever to measurements (i.e., it is an algebraic identity between the various operators involved).

In the asymptotic limit, for n_A and n_C of the form $n_{A/C} \simeq n \pm bn^a$, $n \gg 1$, $a < 1$, the probabilities p_j^n and $p_{j'}^n$ are peaked at $j \simeq rn_A/2$ and $j' \simeq$

$rn_C/2$, as was explained above. Hence, only the average state components $\sigma_{0/1,\xi}^n$ with $\xi = \{j, j'\}$ such that $j \simeq (r/2)n(1 + bn^{a-1})$ and $j' \simeq (r/2)n(1 - bn^{a-1})$ are important. From Eqs. (38) and (39) it is straightforward to obtain

$$\sigma_{0,\xi}^n - \sigma_{1,\xi}^n \simeq r \left(1 - \frac{1-r}{nr^2} \right) (\sigma_0^{rn} - \sigma_1^{rn}) + o(n^{-1}),$$

where we have used that [18] $\langle \hat{J}_z \rangle_j \simeq j - (1-r)/(2r)$ up to exponentially vanishing terms. This relation, for the particular value of $a = 1/2$, is used in the proof of robustness, Eq. (34).

Calculation of Γ_\uparrow

Here we calculate $\Gamma_{\uparrow,\xi} = \text{tr}_B \{ [\uparrow] (\sigma_{0,\xi}^n - \sigma_{1,\xi}^n) \}$, where the average states are defined in Eqs. (27) and (28), and explicitly given in Eqs. (38) and (39) for $\xi = \{j, j'\}$. Let us first calculate the conditional state $\text{tr}_B([\uparrow] \sigma_{0,\xi}^n)$. For that, we need to express $\mathbb{1}_{2j+1}^{AB} = \sum_m [j + \frac{1}{2}, m]$ in the original product basis $\{|j_A, m_A\rangle \otimes |\uparrow/\downarrow\rangle\}$. Recalling the Clebsch-Gordan coefficients $|\langle \frac{1}{2}, \frac{1}{2}; j, m | j + \frac{1}{2}, m + \frac{1}{2} \rangle|^2 = (j+m+1)/(2j+1)$, one readily obtains

$$\text{tr}_B \left([\uparrow] \frac{\mathbb{1}_{2j+1}^{AB}}{d_{2j+1}} \right) = \sum_{m=-j}^j \frac{j+1+m}{2(j+1)d_{2j}} [j, m]_A,$$

which can be written as

$$\text{tr}_B \left([\uparrow] \frac{\mathbb{1}_{2j+1}^{AB}}{d_{2j+1}} \right) = \frac{1}{2} \left(\frac{\mathbb{1}_{2j}^A}{d_{2j}} + \frac{1}{d_{2j}} \frac{\hat{J}_z^A}{j+1} \right),$$

where \hat{J}_z^A is the z component of the total angular momentum operator acting on subsystem A . An analogous expression is obtained for $\text{tr}_B([\uparrow] \mathbb{1}_{2j'+1}^{BC})$. Substituting in Eqs. (38) and (39) and subtracting the resulting expressions, one has $\Gamma_\uparrow = \sum_\xi p_\xi^n \Gamma_{\uparrow,\xi}$, with

$$\Gamma_{\uparrow,\xi} = \frac{1}{2d_{2j_A}d_{2j_C}} \left(\frac{r\langle \hat{J}_z \rangle_{j_A}}{j_A} \frac{\hat{J}_z^A}{j_A+1} - \frac{r\langle \hat{J}_z \rangle_{j_C}}{j_C} \frac{\hat{J}_z^C}{j_C+1} \right), \quad (40)$$

where we have written $\xi = \{j_A, j_C\}$, instead of $\xi = \{j, j'\}$ used in the derivation. For pure states, $r = 1$, $j_A = j_C = n/2$, $\langle \hat{J}_z \rangle_{n/2} = n/2$, and we recover Eq. (11).

In order to minimize the excess risk using SDP, we find it convenient to write Eq. (9) in the form

$$\Delta^{\text{LM}} = 2 \max_{\{\Omega_{m,\xi}\}} \sum_\xi p_\xi^n \text{tr}(\Gamma_{\uparrow,\xi} \Omega_{m,\xi}), \quad (41)$$

where we recall that $m = m_{AC} = m_A + m_C$, and we assumed w.l.o.g. that the seed of the optimal POVM has the block form $\Omega_m = \sum_\xi \Omega_{m,\xi}$. The POVM condition, Eq. (8) must now hold on each block, thus for $\xi = \{j_A, j_C\}$, we must impose that

$$\sum_{m=-j}^j \langle j, m | \Omega_{m,\xi} | j, m \rangle = 2j+1, \quad |j_A - j_C| \leq j \leq j_A + j_C. \quad (42)$$

Programmable machine for unbalanced training sets

The minimum error probability of the optimal programmable machine with a number of n_A , 1 and n_C copies ($n_A > n_C$) in ports A , B and C respectively, is [11]

$$P_e^{\text{opt}} = \frac{1}{4} \left\{ 1 + \frac{D_0}{D_1} - \frac{D_0 + D_1}{D_0 D_1} \sum_{k=0}^{n_C} (n_A - n_C + 2k + 2) \times \sqrt{1 - 4 \frac{D_0 D_1}{(D_0 + D_1)^2} \frac{(n_A - n_C + k + 1)(k + 1)}{(n_A + 1)(n_C + 1)}} \right\},$$

where $D_0 = (n_A + 2)(n_C + 1)$ and $D_1 = (n_A + 1)(n_C + 2)$ are the dimensions of the average states $\sigma_{0/1}$. The asymptotic form of this expression when n_A and n_C are both very large can be easily derived using Euler-Maclaurin's summation formula. The result up to sub-leading order is

$$P_e^{\text{opt}} \simeq \frac{1}{6} \left(1 + \frac{1}{n_A} + \frac{1}{n_C} \right),$$

which leads to Eq. (35).

ACKNOWLEDGEMENTS

We acknowledge financial support from: ERDF: European Regional Development Fund; the Spanish MICINN, through contract FIS2008-01236, FPI Grant No. BES-2009-028117 (GS) and (EB) PR2010-0367; and from the Generalitat de Catalunya CIRIT, contract 2009SGR-0985.

[1] M. Guta and W. Kotlowski, New J. Phys. **12**, 123032 (2010).

[2] C.M. Bishop, *Pattern Recognition and Machine Learning* (Springer, Berlin, 2006).

- [3] Servedio, R. A. & Gortler, S. J., SIAM J. Comput., 33:1067, 2004.
- [4] E. Aïmeur, G. Brassard and S. Gambs, Proc. 19th Canadian Conference on Artificial Intelligence (Canadian AI'06), (Springer-Verlag, Berlin, 2006), pp. 433-444.
- [5] Pudenz, K. L. & Lidar, D. A., arXiv:1109.0325 (2011).
- [6] Bisio, A., Chiribella, G., D'Ariano, G. M., Facchini, S. & Perinotti, P., Phys. Rev. A **81**, 032324 (2010).
- [7] A. Hentschel and B.C. Sanders, Phys. Rev. Lett. **104**, 063603 (2010)
- [8] M. Dusek and V. Buzek, Phys. Rev. A **66**, 022112 (2002); J.A. Bergou and M. Hillery, Phys. Rev. Lett. **94**, 160501 (2005).
- [9] A. Hayashi, M. Horibe, and T. Hashimoto, Phys. Rev. A **72**, 052306 (2005); D. Akimoto and M. Hayashi, Phys. Rev. A **83**, 052328 (2011).
- [10] B. He and J.A. Bergou, Phys. Rev. A **75**, 032316 (2007).
- [11] G. Sentís, *et al.*, Phys. Rev. A **82**, 042312 (2010); **83**, 039909(E), (2011).
- [12] C.W. Helstrom, *Quantum Detection and Estimation Theory*. (Academic Press, New York, 1976).
- [13] E. Bagan, M. Baig and R. Muñoz-Tapia, Phys. Rev. Lett. **87**, 257903 (2001).
- [14] A.S. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (North-Holland, Amsterdam, 1982), p. 163, and references therein.
- [15] R. Derka, V. Buzek and A. K. Ekert, Phys. Rev. Lett. **80**, 1571 (1998).
- [16] A. R. Edmonds, *Angular Momentum in Quantum Mechanics* (Princeton University Press, Princeton, New Jersey, 1960).
- [17] L. Vandenberghe and S. Boyd, SIAM Rev. **38**, 49 (1996).
- [18] B. Gendra, E. Ronco-Bonvehi, J. Calsamiglia, R. Muñoz-Tapia and E. Bagan, arXiv:1205.5479 (2012). Accepted in N. J. Phys.
- [19] A.S. Holevo, Probl. Pereda. Inf. 9 (1973) 3, [Probl. Inf. Transm. 9 (1973) 110].